arXiv:1701.03267v1 [stat.CO] 12 Jan 2017

Robust clustering for functional data based on trimming and constraints

Rivera-García D.[a], García-Escudero L.A.[b], Mayo-Iscar A.[b] and Ortega, J.[a]

[a]CIMAT, A.C. Jalisco s/n, Mineral de Valenciana. Guanajuato 36240, Mexico.
[b]Dept. de Estadística e Investigación Operativa, Universidad de Valladolid. Paseo de Belén, 7. 47005 Valladolid. Spain.

Abstract

Many clustering algorithms when the data are curves or functions have been recently proposed. However, the presence of contamination in the sample of curves can influence the performance of most of them. In this work we propose a robust, model-based clustering method based on an approximation to the "density function" for functional data. The robustness results from the joint application of trimming, for reducing the effect of contaminated observations, and constraints on the variances, for avoiding spurious clusters in the solution. The proposed method has been evaluated through a simulation study. Finally, an application to a real data problem is given.

Keyworks: Functional data analysis · clustering · robustness · functional principal components analysis.

1 Introduction

Recent technological advances have provided more precise instruments, which make possible the recording of large numbers of subsequent measurements in such a way that data can be considered as realizations of random continuous functions. In this context, Functional Data Analysis (Ramsay and Silverman, 2005; Ferraty and Vieu, 2006) has received increasing attention in recent years. Cluster analysis consists of identifying homogeneous groups within a data set and there is also a need for appropriate clustering methods for this new type of functional data sets.

There are many methods to perform cluster analysis for traditional multivariate data among which stand several based on probabilistic models (model-based clustering). The use of the EM algorithm is quite common in order to solve the likelihood maximization involved in all these approaches (Fraley and Raftery, 2002). Posterior probabilities are used to estimate the probabilities of membership of an observation to a specific group.

Several clustering methods for functional data have been recently proposed. A first approximation is known as raw-data clustering, which consists of using the discretization of the functions and directly applying traditional multivariate clustering techniques. A second approximation is based

on a reconstruction of the functional form of the data through the use of basis of functions such as B-splines, Wavelets, Fourier series, etc. (Ramsay and Silverman, 2005). in this case, usual clustering techniques are applied to the fitted coefficients for the functional representation of each curve. Another approach is based on probabilistic models, where a probability distribution for those coefficients is assumed as, for instance in James and Sugar (2003) and more recently Jacques and Preda (2013), who employ an approximation to the "density function" for functional data proposed in (Delaigle and Hall, 2010).

However, the determination of an appropriate clustering technique is even more difficult under the possible presence of outlying curves. One possibility to robustify clustering algorithms is through the application of trimming tools (Cuesta-Albertos et al. (1997), Gallegos (2002), Gallegos (2002)). In García-Escudero et al. (2008), García-Escudero et al. (2014), García-Escudero et al. (2015) and Fritz et al. (2013), restrictions on the matrices of dispersion of the groups are also introduced to avoid the detection of spurious clusters.

Trimming techniques have been already applied as a robust functional clustering tool (García-Escudero and Gordaliza (2005) and Cuesta-Albertos and Fraiman (2007)). This work provides an extension of these principles but in a more model-based approach.

The outline for the rest of this work is as follows. In Section 2, we give a brief description of the approximation to the "density" for functional data that will be applied later. A model-based clustering for functional data is presented in Section 3. Our proposal for robust functional clustering (RFC) and a feasible algorithm for it are described in Section 4. Sections 5 and 6 present a simulation study and real data example to illustrate the performance of the proposed methodology. Finally, we give our conclusions in Section 7.

2 Approximation to the "density function" for functional data

Let $L^2([0,T])$ be a Hilbert space of functions with inner product given by $\langle f,g\rangle = \int f(t)g(t)\,dt$ and norm $||\cdot|| = \langle\cdot,\cdot\rangle^{1/2}$. Suppose X is a random function in $L^2([0,T])$. Assume the process X has mean $\mu(t) = E\{X(t)\}$ and covariance $\Gamma(s,t) = \text{cov}\{X(s),X(t)\}$ which are smooth continuous functions. Consider the Karhunen-Loève (K-L) expansion:

$$X(t) = \mu(t) + \sum_{j=1}^{\infty} C_j(X)\psi_j(t)$$

where the eigenfunctions ψ_j form an orthonormal system and are associated with the covariance operator Γ by means of the corresponding eigenvalues λ_j so that $\langle\Gamma(\cdot,t),\psi_j\rangle = \lambda_j\psi_j(t)$. The eigenfunctions are orthogonal, i.e. they satisfy $\langle\psi_l,\psi_k\rangle = \delta_{lk}$, where δ_{lk} is the Kronecker delta with 1 if $l=k$ and 0 otherwise. The eigenvalues are assumed to be in decreasing order, $\lambda_1 \geq \lambda_2 \geq \cdots$, with $\sum_{j=1}^{\infty}\lambda_j < \infty$. The coefficients $C_j(X)$, $j=1,2,...$, better known as principal components or scores of the stochastic process X, are uncorrelated random variables with zero mean and variance λ_j, such that $C_j(X) = \langle X-\mu,\psi_j\rangle$ is the projection of $X-\mu$ on the j-th eigenfunction ψ_j.

Let $X^{(p)}$ be the approximation of X based on the p first terms in the K-L expansion, this is

$$X^{(p)}(t) = \mu(t) + \sum_{j=1}^{p} C_j(X)\psi_j(t). \tag{1}$$

It can be seen that $E(||X - X^{(p)}||^2) = \sum_{j\geq p+1} \lambda_j$ and $||X - X^{(p)}|| \xrightarrow{m.s.} 0$ when $p \to \infty$.

The notion of density for functional data is not well defined, but there exist some approximations to the density function in the literature. For example, Ferraty and Vieu (2006) have developed extensions of the multivariate case in the nonparametric context.

Without loss of generality, let us suppose that X is a zero mean stochastic process, i.e. $\mu(t) = 0$ for every $t \in [0, T]$ which can be approximated by $X^{(p)}$ as in (1). According to Delaigle and Hall (2010), it is possible to approximate the "density function" for functional data X, when functions are considered in the space determined by the eigenfunctions of the principal components. The notion of small ball probability has an important role in the development of the approximation. Based on the K-L expansion $X^{(p)}$, Delaigle and Hall (2010) show that the probability that X belongs to a ball of radius h centred in $x \in L_2[0, T]$ can be written as

$$\log P(||X - x|| \leq h) = \sum_{j=1}^{p} \log f_{C_j}(c_j(x)) + \xi(h, \rho(h)) + o(\rho(h))$$

where $||X - x||$ denotes the L_2 distance between X and x, f_{C_j} corresponds to the probability density of C_j and $c_j(x) = \langle x, \psi_j \rangle_{L_2}$ is the j-th principal component or score of x. ρ and ξ are functions such that ρ increases to infinity when h decreases to zero. Note that $\log P(||X - x|| \leq h)$ depends on x through the term $\sum_{j=1}^{p} \log f_{C_j}(c_j(x))$. This term captures the first-order effect that x has on $\log P(||X - x|| \leq h)$. Therefore, it serves to describe the main differences in sizes of small-ball probabilities for different values of x since the notion of probability density in the finite dimensional case can be seen as the limit of $P(||X - x|| \leq h)/h$ when h tends to zero. Moreover, as seen in Jacques and Preda (2013), it can be observed that for every $h > 0$ and $x \in L_2([0, T])$,

$$P(||X^{(p)} - x|| \leq h - ||X - X^{(p)}||) \leq P(||X - x|| \leq h) \leq P(||X^{(p)} - x|| \leq h + ||X - X^{(p)}||). \tag{2}$$

Hence, the probability $P(||X - x|| \leq h)$ can be approximated by $P(||X^{(p)} - x|| \leq h)$. If $f_X^{(p)}$ is the joint density function of $C^{(p)} = (C_1, ..., C_p)$ and $x = \sum_{j\geq 1} c_j(x)\psi_j$ then

$$P(||X^{(p)} - x|| \leq h) = \int_{\mathcal{D}_x^{(p)}} f_X^{(p)}(y)\, dy,$$

with $x^{(p)} = \sum_{j=1}^{p} c_j(x)\psi_j$ and $\mathcal{D}_x^{(p)} = \left\{ y \in \mathbb{R}^p : ||y - x^{(p)}||_{\mathbb{R}^p} \leq \sqrt{h^2 - \sum_{j\geq p+1} c_j^2(x)} \right\}$. In this way, the density of $f_X^{(p)}$ can be seen as an approximation of the density of X.

Finally, we can also take into account that the principal components C_j are independent Gaussian random variables when X corresponds to a Gaussian process. In this particular case, $f_X^{(p)}$ is given by

$$f_X^{(p)}(x) = \prod_{j=1}^{p} f_{C_j}(c_j(x))$$

where f_{C_j} is a Gaussian density function with zero mean and variance λ_j.

3 Model-based clustering for functional data

In a clustering framework, we will consider K different models, one for each group. Conditional on the group g, let us consider K-L expansions with p terms and the density function approximation as reviewed in Section 2, such that the density of the principal components in groups are assumed uncorrelated Gaussian variables with zero mean. Moreover, in order to simplify the largely parameterized problem appearing in these clustering frameworks, we consider that the first q_g terms have no restrictions, as in Jacques and Preda (2013), while the remaining $p - q_g$ are constrained in such a way that their corresponding principal components have equal variances, as done in Bouveyron and Jacques (2011). In other words, we assume that scores in each group can be approximated by means of uncorrelated random Gaussian variables with zero mean and covariance matrix $\Sigma_g = \text{diag}(a_{1g}, ..., a_{q_g g}, b_g, ..., b_g)$ with $a_{jg} > b_g$. This means that the main variances of the g-th group are modeled by a_{1g},..., $a_{q_g g}$ while b_g serves to model the variance of the noise of the residual process. If Z_g is a random indicator variable designating membership to group g, for $g = 1, 2, .., K$, then we assume

$$f^{q_g}_{X|Z_g=1}(x) = \prod_{j=1}^{q_g} f_{C_j|Z_g=1}(c_{jg}(x); a_{jg}) \prod_{j=q_g+1}^{p} f_{C_j|Z_g=1}(c_{jg}(x); b_g).$$

Note that when $q_g = p$ we have the model proposed by Jacques and Preda (2013).

Assume now that $Z = (Z_1, ..., Z_K)$ have a multinomial distribution $\mathcal{M}(\pi_1, ..., \pi_K)$, where $\pi_1, ..., \pi_K$ are the mixture probabilities weights (with $\sum_{g=1}^{K} \pi_g = 1$). In this way, the unconditional "approximated functional density" of X at $x \in L^2([0, T])$ is given by

$$f_X(x; \theta) = \sum_{g=1}^{K} \pi_g \left[\prod_{j=1}^{q_g} f_{C_j|Z_g=1}(c_{jg}(x); a_{j,g}) \prod_{j=q_g+1}^{p} f_{C_j|Z_g=1}(c_{jg}(x); b_g) \right],$$

where θ denotes all the parameters that need to be estimated in that expression. Notice that, to start, we are assuming that $(q_1, ..., q_K)$ are known in advance dimensions.

Suppose now that $\{x_1, ..., x_n\}$ is a set of curves being the realization from an independent, identically distributed (i.i.d.) sample from $X \in \mathcal{L}_2([0, T])$. We define the mixture-loglikelihood function as

$$l^p(\theta; x_1, ..., x_n) = \sum_{i=1}^{n} \log \left(\sum_{g=1}^{K} \pi_g \left[\prod_{j=1}^{q_g} \frac{1}{\sqrt{2\pi a_{jg}}} \exp\left(\frac{-c^2_{ijg}}{2a_{jg}} \right) \prod_{j=q_g+1}^{p} \frac{1}{\sqrt{2\pi b_g}} \exp\left(\frac{-c^2_{ijg}}{2b_g} \right) \right] \right),$$

where $c_{ijg} = c_{jg}(x_i)$ corresponds to the j-th principal component of the curve x_i in group g.

4 Robust functional clustering based on trimming and constraints

García-Escudero et al. (2014) and Ritter (2015) provide a methodology for robust mixture modeling in a multivariate real-valued context. This methodology proposes using trimming and scatter constraints to remove the contamination in the data and simultaneously avoid spurious clusters. The

methodology is also based on a particular type of trimming which is determined by the dataset itself. To be more precise, if $\{x_1, ..., x_n\}$ is a random sample in $\mathbb{R}^p$, the idea is to maximize the trimmed mixture likelihood defined as

$$\sum_{i=1}^{n} \eta(x_i) \log \left[\sum_{g=1}^{K} \pi_g \phi(x_i; \theta_g) \right], \tag{3}$$

where $\phi(\cdot; \theta_g)$ stands for the p-dimensional Gaussian density with parameters $\theta_g = (\mu_g, \Sigma_g)$. The indicator function $\eta(\cdot)$ serves to designate whether the observation x_i has been trimmed (when $\eta(x_i) = 0$) or not (when $\eta(x_i) = 1$). A proportion α of observations is trimmed, so that $\sum_{i=1}^{n} \eta(x_i) = [n(1-\alpha)]$. Constraints on the eigenvalues of the scatter matrices are also applied in order to make the maximization problem well defined and to avoid the detection of non-interesting spurious solutions (García-Escudero et al. (2008)). Then, it is proposed to maximize (3) subject to the restriction

$$\frac{\max_{g,j} \lambda_j(\Sigma_g)}{\min_{g,j} \lambda_j(\Sigma_g)} \leq d,$$

where $\{\lambda_j(\Sigma_g)\}_{j=1}^{p}$ is the eigenvalue set for matrix Σ_g and $d \geq 1$ is a fixed constant. In the most constrained case ($d = 1$), we are searching for homoscedastic and spherical clusters.

In a similar fashion, we can adapt this methodology for functional data by considering a trimmed and constrained version of the model-based clustering approach presented in Section 3. Let $\{x_1, ..., x_n\}$ be a realization from a i.i.d. sample of the process $X \in L^2[0, T]$. A trimmed loglikelihood can be defined in this functional setting as

$$l_\alpha^p(\theta; x_1, ..., x_n) = \sum_{i=1}^{n} \eta(x_i) \log \left(\sum_{g=1}^{K} \pi_g \left[\prod_{j=1}^{q_g} \frac{1}{\sqrt{2\pi a_{jg}}} \exp\left(\frac{-c_{ijg}^2}{2a_{jg}} \right) \prod_{j=q_g+1}^{p} \frac{1}{\sqrt{2\pi b_g}} \exp\left(\frac{-c_{ijg}^2}{2b_g} \right) \right] \right) \tag{4}$$

where $c_{ijg} = c_{jg}(x_i)$ is the j-th principal component corresponding to curve x_i in group g and, again, $\sum_{i=1}^{n} \eta(x_i) = [n(1-\alpha)]$. To avoid spurious solutions, we set two constants d_1 and d_2, both greater or equal than 1, and impose the following constraints on the scatter parameters:

$$\frac{\max_{g=1,...,K; j=1,...,q_j} a_{jg}}{\min_{g=1,...,K; j=1,...,q_j} a_{jg}} \leq d_1$$

and

$$\frac{\max_{g=1,...,K} b_g}{\min_{g=1,...,K} b_g} \leq d_2.$$

4.1 Proposed algorithm

Of course, the maximization of the trimmed log-likelihood in (4) may not be an easy task from a computational point of view. A classical way of maximizing mixture model likelihoods is to use the EM algorithm. The algorithm proposed here is based on the traditional EM algorithm incorporating some additional steps. In a so-called T-step (Trimming step) we temporally discard those observations

with smallest contributions to the likelihood (to increase as much as possible the trimmed log-likelihood). We also consider, in the M-step, a final refinement where the required constraints on the scatter matrices are imposed on the scatter parameters.

The proposed algorithm may be described as follows, where $\theta^{(l)}$ are the values of parameters at stage l of the iterative process:

1. *Initialization:* The algorithm is randomly initialized `nstart` times by selecting different starting $\theta^{(0)}$ parameters. With this idea in mind, we simply propose to randomly select $K \times h$ subindexes $\{i_{g_1}, i_{g_2}, ..., i_{g_h}\}_{g=1}^{K} \subset \{1, 2, ..., n\}$ where h is the minimum number of observations needed to computationally carry out a functional principal component analysis for these observations. We then apply the procedure that will be latter described in Step 2 of this algorithm with weights $\tau_{i_{g_1} g} = \tau_{i_{g_2} g} = ... = \tau_{i_{g_h} g} = 1$, $g = 1, ..., K$, and weights $\tau_{ig} = 0$ for all the remaining (i, g) pairs. The smaller the h the more likely is that these $K \times h$ observations could be free of outliers (or at least with not so many within) in any of those random initializations.

2. *Trimmed EM steps:* The following steps are alternatively executed until convergence (i.e. $\theta^{(l+1)} = \theta^{(l)}$) or a maximum number of iterations `iter.max` is reached.

   2.1. *T- and E-steps:* Let us use the notation

$$D_g(x_i, \theta) = \pi_g \prod_{j=1}^{q_g} \frac{1}{\sqrt{2\pi a_{jg}}} \exp\left(\frac{-c_{ijg}^2}{2a_{jg}}\right) \prod_{j=q_g+1}^{p} \frac{1}{\sqrt{2\pi b_g}} \exp\left(\frac{-c_{ijg}^2}{2b_g}\right)$$

   and

$$D(x_i, \theta) = \sum_{g=1}^{K} D_g(x_i, \theta).$$

   If we consider $D(x_{(1)}; \theta^{(l)}) \leq D(x_{(2)}; \theta^{(l)}) \leq .... \leq D(x_{(n)}; \theta^{(l)})$, the observations with indexes in

$$I = \{i : D(x_i; \theta^{(l)}) \leq D(x_{([n\alpha])}; \theta^{(l)})\} \tag{5}$$

   are those which are tentatively discarded in this iteration of the algorithm.

   As in other mixture fitting EM algorithms, we compute posterior probabilities by using the well-known Bayes rule as

$$\tau_g(x_i; \theta^{(l)}) = D_g(x_i; \theta^{(l)}) / D(x_i; \theta^{(l)}), \text{ for } i = 1, ..., n.$$

   However, unlike standard EM algorithms, the $\tau_g(x_i; \theta^{(l)})$ values for the discarded observations are modified as

$$\tau_g(x_i; \theta^{(l)}) = 0, \text{ for all } g = 1, ..., K, \text{ when } i \in I.$$

   Notice that the way that trimming is done is similar to that in García-Escudero et al. (2014).

   2.2. *M-step:* This step consists of three stages:

2.2.1 *Weights update:* Weights are updated as

$$\pi_g^{(l+1)} = \sum_{i=1}^{n} \tau_g(x_i; \theta^{(l)})/[n(1-\alpha)]$$

2.2.2 *Principal component update:* Consider a basis of functions $\Phi = \{\phi_1, ..., \phi_p\}$. If x_i admits an approximate reconstruction in this basis as $x_i(t) \simeq \sum_{j=1}^{p} \gamma_{ij}\phi_j(t)$ then let Γ be the $n \times p$ matrix of coefficients γ_{ij} used in that reconstruction. Let W be the matrix of the inner products between the basis functions $W_{jl} = \int_0^T \phi_j(t)\phi_l(t)dt$ $(1 \leq j, l \leq p)$. The updating of the principal components is carried out by weighting the importance of the untrimmed $x_i(t)$ curves by the conditional probability $T_g^{(l)} = \text{diag}(\tau_{1,g}^{(l)}, ..., \tau_{n,g}^{(l)})$. The first step is to center the curve $x_i(t)$ in group g, by subtracting the weighted pointwise sample mean calculated with $\tau_{i,g}^{(l)}$ weights. The expansion coefficients of the centered curves are given by $\Gamma_g^{(l)} = (I_n - 1_n(\tau_{1,g}^{(l)}, ..., \tau_{n,g}^{(l)}))\Gamma$, where I_n is the $n \times n$ identity matrix and $1_n = (1, 1, ..., 1)$ is the unit vector. Note that the weighted sample covariance function is then given by

$$v^{(l+1)}(s,t) = \frac{1}{n_g^{(l)}} \sum_{i=1}^{n} \tau_{ig}^{(l)} x_i(s)x_i(t), \tag{6}$$

where $n_g^{(l)} = \sum_{i=1}^{n} \tau_{ig}^{(l)}$. Consider also that the j-th eigenfunction can be written as $\psi_j(s) = \beta_j^T \phi(s)$ with $\phi(t) = (\phi_1(t), ..., \phi_p(t))'$. Substituting the above expressions into (6) one obtains

$$v^{(l+1)}(s,t) = (n_g^{(l)})^{-1}\phi^T(s)\Gamma_g^{T(l)}T_g^{(l)}\Gamma_g^{(l)}\phi(t),$$

and one gets the eigenequation

$$(n_g^{(l)})^{-1}\phi^T(s)\Gamma_g^T T_g^{(l)}\Gamma_g \left(\int \phi(t)\phi^T(t)dt\right)\beta_j = \lambda\phi^T(s)\beta_j.$$

By using the W matrix of the inner products, the previous equation can be written as

$$(n_g^{(l)})^{-1}\phi^T(s)\Gamma_g^{T(l)}T_g^{(l)}\Gamma_g^{(l)}W\beta_j = \lambda\phi^T(s)\beta_j.$$

Observing that the previous expression is valid for all values of s, one gets

$$(n_g^{(l)})^{-1}\Gamma_g^{T(l)}T_g^{(l)}\Gamma_g^{(l)}W\beta_j = \lambda\beta_j,$$

with the additional constraint $\|\psi_j\|^2 = 1$ that turns into $\beta^T W\beta = 1$. Let us define $u_j = W^{1/2}\beta_j$ and, then, the following eigenequation is finally obtained:

$$(n_g^{(l)})^{-1}W^{1/2}\Gamma_g^{T(l)}T_g^{(l)}\Gamma_g^{(l)}W^{1/2}u_j = \lambda u_j$$

subject to $u^T u_j = \|u_j\|^2 = 1$. From this equation we can compute the eigenvalues λ_j and the vector of coefficients $\beta_j = W^{-1/2}u_j$, with which we calculate the eigenfunctions $\psi_j(s)$ and the principal component scores $C_{jg}^{(l+1)}$ are given by $C_{jg}^{(l+1)} = \Gamma_g^{(l)}W\beta_j$.

2.2.3 *Scatter parameters update:* The parameters $a_{1,g}, ..., a_{q_g,g}$ and b_g are initially estimated as $\hat{a}_{j,g}^{(l+1)} = \lambda_j$ for the q_g first eigenvalues of the $W^{1/2}\Gamma_g^{T(l)}T_g^{(l)}\Gamma_g^{(l)}W^{1/2}$ matrix and

$$b_g^{(l+1)} = \frac{1}{p - q_g}\left[\text{trace}\left(W^{1/2}\Gamma_g^{T(l)}T_g^{(l)}\Gamma_g^{(l)}W^{1/2}\right) - \sum_{j=1}^{q_g}\hat{a}_{jg}^{(l+1)}\right].$$

Recall that the scatter parameters so obtained do not necessarily satisfy the required constraints for the given d_1 and d_2 constants. In case that these constraints do not hold, following Fritz et al. (2013), we define their truncated versions as:

$$a_{jg}^{m_1} = \begin{cases} a_{jg} & \text{if } a_{jg} \in [m_1, d_1 m_1], \\ m_1 & \text{if } a_{jg} < m_1, \\ d_1 m_1 & \text{if } a_{jg} > d_1 m_1, \end{cases}$$

and

$$b_g^{m_2} = \begin{cases} b_g & \text{if } b_g \in [m_2, d_2 m_2], \\ m_2 & \text{if } b_g < m_2, \\ d_2 m_2 & \text{if } b_g > d_2 m_2. \end{cases}$$

The scatter parameters are finally updated as $\{a_{1,g}^{m_{opt_1}}, ..., a_{q_g,g}^{m_{opt_1}}, b_g^{m_{opt_2}}, ..., b_g^{m_{opt_2}}\}$ where m_{opt_1} minimizes

$$m_1 \mapsto \sum_{g=1}^{K} n_g \sum_{j=1}^{q_g}\left(\log(a_{jg}^{m_1}) + \frac{a_{jg}}{a_{jg}^{m_1}}\right),$$

and, m_{opt_2} minimizes

$$m_2 \mapsto \sum_{g=1}^{K} n_g (p - q_g)\left(\log(b_g^{m_2}) + \frac{b_g}{b_g^{m_2}}\right)$$

where $n_g = \sum_{i=1}^{n} \tau_{ig}$. These are indeed two real-valued functions that can be easily minimized (see Fritz et al. (2013)).

3. *Evaluate target function:* After applying the trimmed EM steps, the associated value of the target function (4) is computed (we set $\eta(x_i) = 0$ if $i \in I$ and $\eta(x_i) = 1$ if $i \notin I$ for I defined as in (5) with the final iteration parameters). The set of parameters yielding the highest value of this target function and the associated trimmed indicator function η are returned as the final algorithm's output.

4.2 Estimation of dimensions

As in Bouveyron and Jacques (2011) and Jacques and Preda (2013), the estimation of the dimensions per group, q_g, $g = 1, ..., K$ in the K-L expansion is not an easy task and still an open problem. In the previously mentioned works, the authors used the "Cattell" procedure (Cattell, 1966) and show that, by using an appropriate threshold, K sensible values can be obtained. However, the application of the "Cattell" threshold within the EM algorithm may create increments and decrements of the

target function between two successive iterations. In this work, we prefer solving the maximization of the target function for fixed combinations of dimension and, later, choose the dimensions yielding the better value of a penalized likelihood for fixed values of trimming levels α and constraints d_1 and d_2. To be more precise, we choose the dimensions minimizing the Bayesian Information Criterion (BIC) defined as

$$BIC = -2l_{\alpha}^{p}(\hat{\theta}; x_1, ..., x_n) + \kappa \log(n)$$

where $l_{\alpha}^{p}(\hat{\theta})$ corresponds to the trimmed log-likelihood function valued at the estimated optimal parameters $\hat{\theta}$, n is the number of observations and κ corresponds to the number of free parameters to be estimated. We have $\kappa = \rho + \nu + 2K + Q$, where $\rho = (Kp + K - 1)$ is the number of parameters needed to estimate means and mixture proportions, $\nu = \sum_{g=1}^{K} q_g[p - (q_g + 1)/2]$ corresponds to the number of parameters needed to estimate the ψ_j eigenfunctions and $Q = \sum_{g=1}^{K} q_g$.

To illustrate the use of BIC in the selection of the dimensions q in the K-L expansion, we simulate a data set from the simulation scheme called Scenario 1 ($q_1 = 2$ and $q_2 = 3$) with 10% contamination of type iii) as will be fully described in the Section 5). Figure 1 shows that simulated dataset. Figure 2 (right panel) shows the BIC values for several dimension combinations when $d_1 = d_2 = 10$ and $\alpha = 0.1$. After, testing those combinations, it was observed that the minimum value for the BIC corresponds to dimension $q_1 = 2$ and $q_2 = 3$. Moreover, as can be seen in Figure 2 (left panel), we note that the minimum value of the BIC corresponds to one of the best solutions of the algorithm in terms of correct classification rate (CCR).

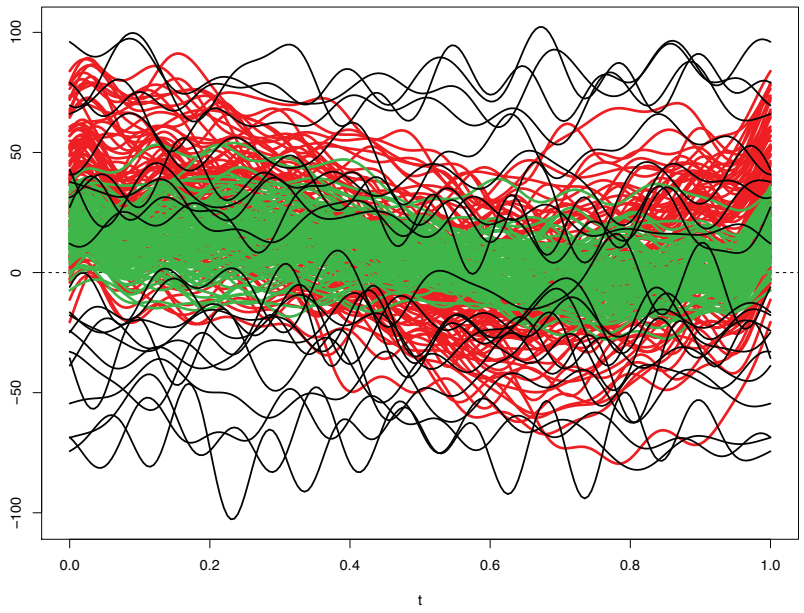

Figure 1: A simulated data set with $K = 2$ groups from "Scenario 1 and contamination scheme (iii)" (as described in Section 5). The subspace dimensions in this example are $q_1 = 2$ and $q_2 = 3$.

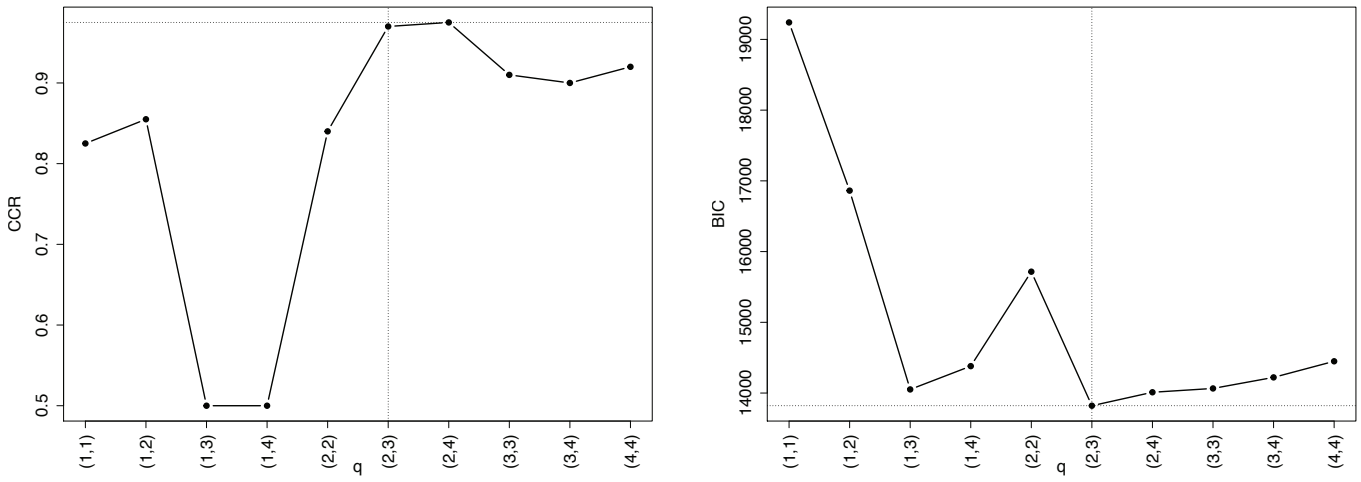

Figure 2: Selection of the dimensions by means of BIC for the data set in Figure 1 when $d_1 = d_2 = 10$ and $\alpha = 0.1$. Different combinations of dimensions are represented in the x-axis, BIC values are presented in the right panel while the corresponding correct classification rates are in the left panel.

5 Simulation study

In order to evaluate the performance of the methodology proposed, we simulated different scenarios and contamination types.

For the "good" observations arranged in $K = 2$ clusters, we consider the following scheme of simulation based on the K-L expansion:

$$x_i(t) = \mu_g(t) + \sum_{j=1}^{q_g} a_{jg}^{1/2} z_i \psi_j(t) + \sum_{j=q_g+1}^{p} b_g^{1/2} z_i \psi_j(t) \quad t \in [0,1], \tag{7}$$

where z_i are independent and $N(0,1)$-distributed, μ_g are the group mean function, a_{jg} corresponds to the main variances and b_g corresponds to the residual variability. In this simulation, we consider that the eigenfunctions ψ_j are the first 21 Fourier basis functions that are defined as

$$\psi_j(t) = \begin{cases} \psi_0(t) = 1, \\ \psi_{2j-1}(t) = \sqrt{2}\sin(j2\pi t) \\ \psi_{2j}(t) = \sqrt{2}\cos(j2\pi t) \end{cases},$$

for $j = 1, 2, ..., p$. We assume that the first $i = 1, ..., 100$ observations are generated when $g = 1$ in (7) and the second group of observations with indices $i = 101, ..., 200$ are generated when $g = 2$ in (7). We have two different main scenarios for the "good" part of data depending on the mean functions and chosen variances:

Scenario 1: The groups have the same mean $\mu_1(t) = \mu_2(t) = \cos(t)$ and dimensions $q_1 = 2$ and $q_2 = 3$. The variances for the first group are $(a_{11}, a_{21}) = (60, 30)$ and $b_1 = 0.5$. For the second group, the variances are $(a_{12}, a_{22}, a_{32}) = (170, 140, 120)$ and $b_2 = 1$.

Scenario 2: The groups have different means $\mu_1(t) = \cos(t) + 3$ and $\mu_2(t) = \cos(t) + 1$ and the dimensions are $q_1 = 2$ and $q_2 = 3$. The variances are $(a_{11}, a_{21}) = (a_{12}, a_{22}) = (60, 30)$, $b_1 = 0.5$ and $b_2 = 1$.

We also consider the possibility of adding another 22 curves (10% contamination level) to see the effect of noise in clustering. In two out of the three contaminating schemes, each of these 22 contaminating curves x_i are obtained by fitting a linear combination of the 21 first Fourier base elements plus a global mean which interpolates 21 points in $\mathbb{R}^2$ as

$$\{(t_l, u_i + \varepsilon_{i,l})\}_{l=1}^{21},$$

where $\{t_l\}_{l=1}^{21}$ is an equispaced grid on $[0, 1]$, $\{u_i\}_{i=1}^{22}$ is the result of random sample from a uniform distribution in the $[a, b]$ interval (to be specified latter) and $\{\varepsilon_{i,l}\}_{l=1}^{21}$, for $i = 201, ..., 222$, are independent normally distributed error terms with variance $\sigma^2 = 10$.

For both scenarios (Scenario 1 and 2) for the "good" part of data, we consider the following contaminating schemes:

(i) No contamination (i.e., the total number of observations is $n = 200$).

(ii) Using the previously described contaminating scheme with $[a, b] = [150, 180]$. This means that the contaminating curves are clearly far apart from the "good" curves.

(iii) Using the previously described contaminating scheme with

$$[a, b] = \left[\min_{i=1,...,200; t\in[0,1]} x_i(t), \max_{i=1,...,200; t\in[0,1]} x_i(t) \right].$$

(iv) We use the scheme in (7) also for $i = 201, ..., 222$ but the normally distributed z_{ij} variables are replaced by (heavier tailed) Cauchy distributed ones.

In order to test the performance of the methodology proposed here, we carry out a simulation study using the scheme previously described and compare the results with those obtained by "Funclust" (Jacques and Preda, 2013) and "FunHDDC" (Bouveyron and Jacques, 2011).

In this simulation study, it is important to note that we assume the q_g dimensions to be unknown parameters and that we use the BIC proposal described in section 4.2 to estimate them when applying the proposed robust functional clustering (RFC). We use trimming levels $\alpha = 0$ (untrimmed) and $\alpha = 0.1$, constraints $d_1 = d_2 = 1$, $d_1 = d_2 = 10$ and $d_1 = d_2 = 10^{10}$ (i.e., almost unconstrained in this last case). We always return the best solution in terms of the highest BIC value for each combination of all those fixed values of trimming levels and constraints. We use `niter`= 100 random initializations with `iter.max`= 20.

For the "Funclust" method we have used the library `Funclustering` (Soueidatt, 2014) in R where the EM algorithm has been initialized with the best solutions out of 20 "short" EM algorithms with only 20 iterations with values of $\varepsilon = 0.001, 0.05, 0.1$ in the Cattell test. In the case of the "FunHDDC", we use the library `FunHDDC` (Bouveyron and Jacques, 2014) in R with values of $\varepsilon = 0.001, 0.05, 0.1$ in the Cattell test, moreover, the submodels $A_{kj}B_kQ_kD_k$, $A_kB_kQ_kD_k$, $A_kBQ_kD_k$, $AB_kQ_kD_k$, ABQ_kD_k are tested, see details in (Bouveyron and Jacques, 2011) and the best solution in terms of the highest BIC value for all those submodels are returned.

Figure 3 shows the results for datasets simulated according to Scenario 1, i.e. groups with equal means. This figure is composed of a matrix of graphs, where the rows correspond to the different contamination schemes (uncontaminated in the first row) while the columns correspond to the methodologies tested. The first column corresponds to "Funclust", the second to "FunHDCC" and the third one shows the results for the robust functional clustering (RFC) procedure with the two different trimming levels and the three constraints levels (we are assuming $d_1 = d_2$ to simplify the simulation study). The x-axis corresponds to the threshold in the Cattell test for the first two columns, and to the constraint level for RFC, while the y-axis corresponds to the correct classification rate (CCR).

The results show that the joint use of trimming and constraints in RFC improve the CCR substantially. Results are very good for moderate ($d_1 = d_2 = 10$) and small ($d_1 = d_2 = 1$) values of the constraint constants, while for high values the results are poor. Very high values for these constants are equivalent to having unconstrained parameters. The use of trimming also turns out to be very useful in all the contaminated case while it does not affect so much the results in the uncontaminated case.

In most cases the results for "FunHDDC" and "Funclust" fall below those of RFC when applying the $\alpha = 0.1$ trimming and small/moderate values d_1 and d_2 for the variance parameters. The only case where this is not so is "Funclust" with $\tau = 0.001$ in the first row, corresponding to uncontaminated data. However, this method requires the use of $q_1 = 20$ and $q_2 = 8$ terms in the K-L expansion for groups 1 and 2 respectively.

The results corresponding to Scenario 2 are presented in Figure 4. This scenario corresponds to groups with different means. Again, it can be seen that the joint use of trimming and constraints improve the results in terms of classification rates. The results in these cases, both for moderate ($d_1 = d_2 = 10$) and small ($d_1 = d_2 = 1$) values of the constraint constants are quite good, while the results are poor for very large $d_1 = d_2$ values. In this case the RFC method with appropriate trimming and constraints always performs better than "FunHDDC" and "Funclust" in terms of classification accuracy.

In addition, it is worth mentioning that the results for RFC method for both simulation scenarios are more consistent, in the sense that the correct classification rate (CCR) has a lower dispersion for this method, which indicates another advantage of this proposal for robust clustering.

6 Real data example: NOx levels

The data set corresponds to daily curves of Nitrogen Oxides NOx emissions in the neighborhood of the industrial area of Poblenou, Barcelona (Spain). NOx is one of the principal contaminant agents and characterizing its behavior is useful to develop appropriate environmental policies. The detection of outlying emission curves from any data source is meaningful because the explanation of why these curves are observed may be helpful in order to forecast or anticipate them. In addition, these outlying curves can also influence non-robust clustering methods leading to wrong conclusions when searching for clusters of days with different types of emission patterns.

The data are available in the `fda.usc` library (Febrero-Bande and Oviedo de la Fuente, 2012) in R. The measurements of NOx (in $\mu g/m^3$) were taken hourly resulting in 115 days with complete observations. This data set has been analyzed to test methodologies for the detection of outliers in

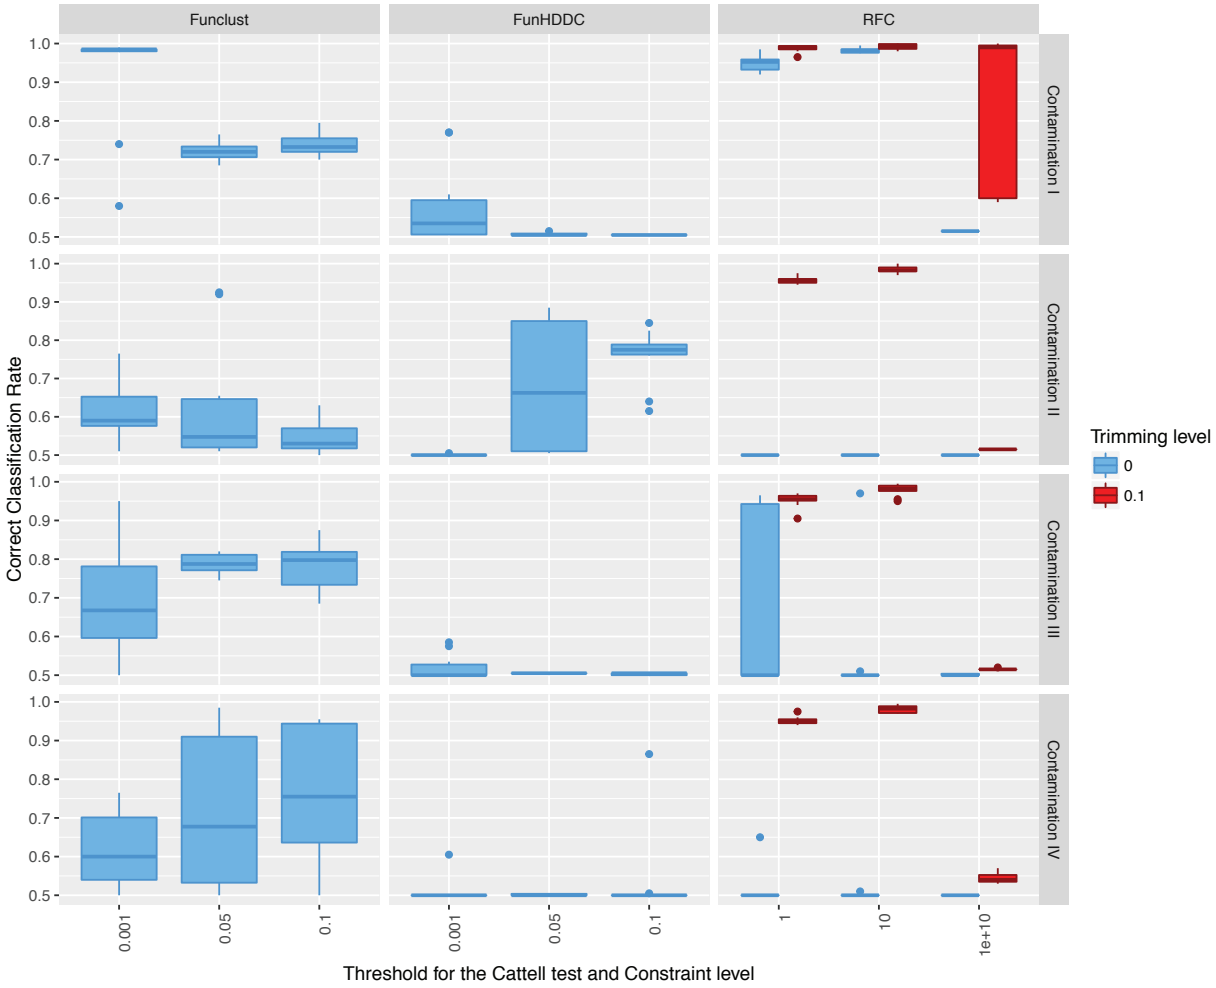

Figure 3: Scenario 1 (equal mean functions): Correct classification rate (CCR) for the three methods considered, represented in different columns. Rows correspond to the different contamination schemes (i) to (iv), described previously in this section, starting with no contamination in the first row. Constraint levels $d_1 = d_2 = 1$, 10 and 10^{10} and trimming levels $\alpha = 0$ and 0.1 are used for the RFC method and the proposed BIC to choose dimension. Threshold values $\varepsilon = 0.001, 0.05$ and 0.1 are used for the "Cattell" procedure in "Funclust" and "FunHDDC".

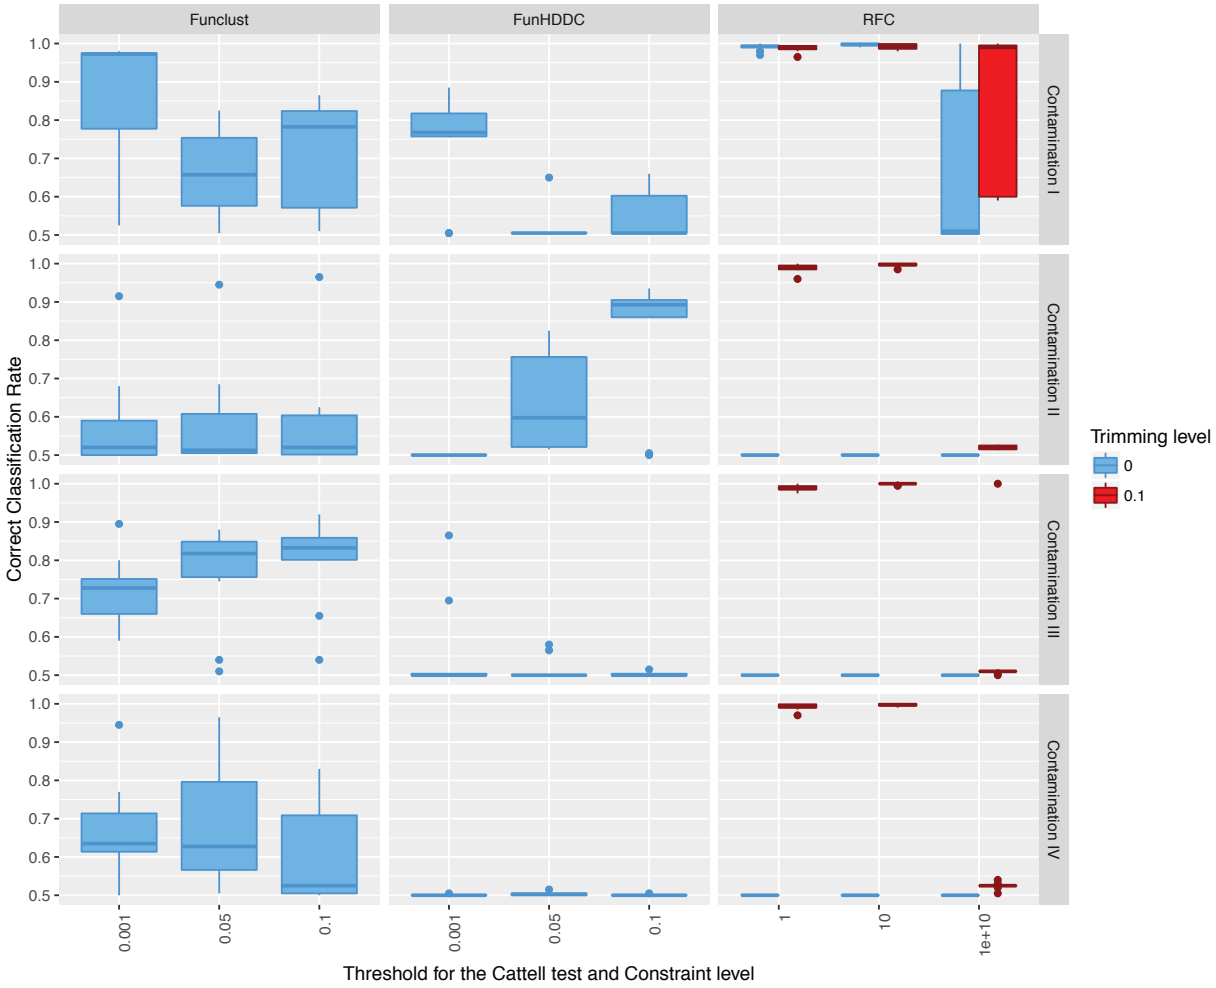

Figure 4: Scenario 2 (unequal mean functions): Correct classification rate (CCR) for the three methods considered, represented in different columns. Rows correspond to the different contamination schemes (i) to (iv), described previously in this section, starting with no contamination in the first row. Constraint levels $d_1 = d_2 = 1$, 10 and 10^{10} and trimming levels $\alpha = 0$ and 0.1 are used for the RFC method and the proposed BIC to choose dimension. Threshold values $\varepsilon = 0.001, 0.05$ and 0.1 are used for the "Cattell" procedure in "Funclust" and "FunHDDC".

functional data in Febrero et al. (2008), Sguera et al. (2015) and Sawant et al. (2012).

Figure 5 shows the original daily curves of Nitrogen Oxides NOx emissions by using different colors, red for the 76 "working days" and green for the 39 "non-working days".

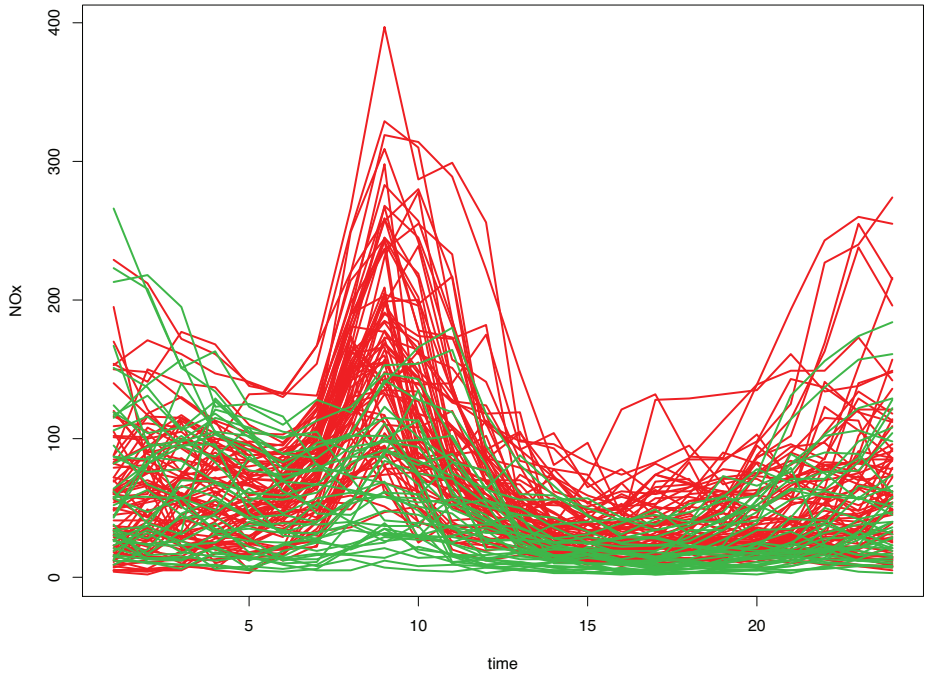

Figure 5: Curves represent daily levels of NOx for 115 days, with 76 working day in red and 39 non-working days in green.

The RFC methodology is applied to this dataset and the results are compared to those obtained using the "Funclust" and "FunHDCC" methodologies. Two clusters ($K = 2$) and a B-spline basis of functions of order 3 with 15 basis elements (13 equispaced knots) are taken. For RFC, we use trimming levels $\alpha = 0$, 0.1 and 0.15, and constraints values $d_1 = d_2$ equal to 1, 10 and 10^{10} (`nstart`= 100 and `niter`= 20). For the "Funclust" and "FunHDCC" methods, we use the same strategy as in the simulation study with values of $\varepsilon = 0.001$, 0.05 and 0.1 in the Cattell procedure. The dimensions are estimated by using the BIC criterion for the RFC method and also when applying the "FunHDDC" method.

Table 1 shows a summary of the results obtained for different combinations of input parameters. The second column of the table shows the estimated q_g dimensions by means of the BIC for RFC method and the Cattell test for "Funclust" and "FunHDDC" methods. We also give the correct classification rates ("CCR" column) assuming that the "true" clusters in data were only determined by the type of day (working and non-working days). In this column, again, we are re-assigning the trimmed observations according to their posterior probabilities of membership to clusters.

One can see that the use of strong $d_1 = d_2 = 1$ constraints slightly increases the CCR (assuming that the correct groups were determined by working and non-working days). In this case, the CCR for RFC is 84.3% without trimming and 85% with a $\alpha = 0.1$ and 84.3% with $\alpha = 0.15$ while the

	q_g	α	d_1	d_2	ε	CCR
RFC	2,5	0	1	1	-	0.84
	5,5	0	10	10	-	0.70
	5,5	0	10^{10}	10^{10}	-	0.69
	2,5	0.1	1	1	-	0.85
	5,5	0.1	10	10	-	0.69
	5,5	0.1	10^{10}	10^{10}	-	0.66
	2,5	0.15	1	1	-	0.84
	5,5	0.15	10	10	-	0.70
	5,5	0.15	10^{10}	10^{10}	-	0.69
Funclust	14,13	-	-	-	0.001	0.84
	4,5	-	-	-	0.05	0.66
	3,3	-	-	-	0.1	0.66
FunHDDC	14,10	-	-	-	0.001	0.66
	3,2	-	-	-	0.05	0.66
	1,3	-	-	-	0.1	0.66

Table 1: Correct classification rate (CCR) and dimension estimated for different levels of trimming α and constraints d_1 and d_2, for the RFC method and different vales of ε for the Cattell test in "Funclust" and "FunHDDC".

best CCR for "Funclust" is 84.3%. However, the RFC method has an additional advantage in that it requires smaller dimensions than "Funclust" for achieving that level of CCR.

Another important point is that the RFC allows us to perform clustering and outlier detection simultaneously while "Funclust" and "FunHDDC" do not. Even though the detected outliers are not so extreme in this case as to completely deteriorate the clustering process, it is also interesting to detect these outlying curves also taking the cluster structure in mind. In this direction, every trimmed curve (trimming levels $\alpha = 0.1$ and $\alpha = 0.15$) corresponds to outliers already detected in previous works in the literature that were also concerned with functional outlying detection as Febrero et al. (2008)(DEPTH), Sguera et al. (2015)(KFSD) and Sawant et al. (2012)(BACONPCA). Two separated data sets, considering only working days (W) and non-working days (NW), were considered when applying Sguera et al. (2015)(KFSD) while the complete dataset, without differentiating between working and non-working days (W-NW), is used when applying our RFC proposal and the other two methods.

Figure 6 shows the RFC clustering results. We observe that the curves that are detected as outliers (in black in the third column) exhibit different patterns from the rest of the curves.

7 Conclusions

A feasible methodology for robust model-based functional clustering has been proposed and illustrated. The key idea behind the algorithm presented is the use of an approximation of the "density" for functional data together with the simultaneous use of trimming and constraints. This allows for a robust model-based clustering approach.

The use of trimming tools protects the estimation of the parameters against the harmful effect

KFSD		BACONPCA	DEPTH	RFC $\alpha = 0.1$	RFC$\alpha = 0.15$
NW	W	W-NW	W-NW	W-NW	W-NW
12/03/2005	09/03/2005	18/03/2005	11/03/2005	25/02/2005	25/02/2005
19/03/2005	11/03/2005	29/04/2005	18/03/2005	03/03/2005	03/03/2005
30/04/2005	15/03/2005	11/03/2005	29/04/2005	11/03/2005	09/03/2005
01/05/2005	16/03/2005	02/05/2005	02/05/2005	16/03/2005	11/03/2005
	17/03/2005	09/03/2005		18/03/2005	16/03/2005
	18/03/2005			25/04/2005	18/03/2005
	29/04/2005			29/04/2005	18/04/2005
	02/05/2005			02/05/2005	25/04/2005
				18/05/2005	29/04/2005
				27/05/2005	02/05/2005
				23/06/2005	03/05/2005
				15/05/2005	18/05/2005
					27/05/2005
					23/06/2005
					19/03/2005
					30/04/2005
					15/05/2005

Table 2: Outliers detected when using Febrero et al. (2008)(DEPTH), Sguera et al. (2015)(KFSD), Sawant et al. (2012)(BACONPCA) and the proposed RFC methodology with $\alpha = 0.1$ and $\alpha = 0.15$. Separated data sets considering only working days (W) and non-working days (NW) were used by Sguera et al. (2015)(KFSD) while the complete data set (W-NW) were used for the other methods.

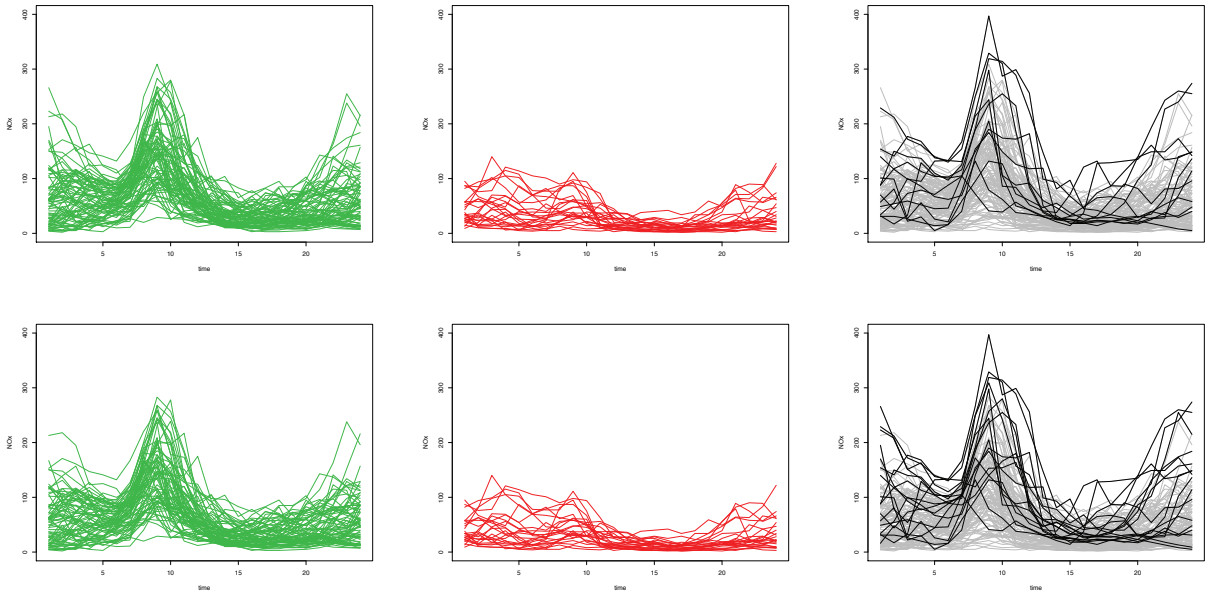

Figure 6: Clusters found (non-trimmed curves in green and red) when applying the RFC method with $K = 2$ and $d_1 = d_2 = 1$. The trimmed curves appear in black while the non-trimmed ones in gray. *Top panels:* Trimming level $\alpha = 0.1$. *Bottom panels:* Trimming level $\alpha = 0.15$.

of (even a small amount of) outlying curves, while the constraints avoid the detection of spurious clusters and improve the algorithm's stability. The simulation study shows that the joint use of constraints and trimming tools improve the performance of the clustering algorithm in comparison to some other procedures for clustering functional data. The real data example shows that the trimmed curves often correspond to outliers already detected by other specialized methods for outlier detection in functional data analysis. In fact, we conclude that the proposed robust methodology can be a useful tool to detected contamination and groups in a functional data set simultaneously.

However, some limitations of this methodology are the choice of level of trimming α and the choice of the scatter constraints constants d_1 and d_2. These values are subjective and sometimes depend on the final purpose of the cluster analysis. For this reason, we always recommend the use of different values of trimming and constraint and monitoring the effect in the clustering partition of these choices. The development of more automatized selection procedures for these values may be considered as an open problem for future research.

Finally, an extension of our proposal for future work is the consideration of multivariate functional data.

8 Acknowledgements

This work was partly done while DR and JO visited the Departamento de Estadística e I.O., Universidad de Valladolid, Spain, with support from Conacyt, Mexico (DR as visiting graduate student, JO by Projects 169175 Análisis Estadístico de Olas Marinas, Fase II y 234057 Análisis Espectral, Datos Funcionales y Aplicaciones), CIMAT, A.C. and the Universidad de Valladolid. Their hospitality and support is gratefully acknowledged. Research by LA G-E and A M-I was partially supported by the Spanish Ministerio de Economía y Competitividad y fondos FEDER, grant MTM2014-56235-C2-1-P, and by Consejería de Educación de la Junta de Castilla y León, grant VA212U13

Bibliography

Bouveyron, C. and J. Jacques (2011). Model-based clustering of time series in group-specific functional subspaces. *Adv. Data Anal. Classif. 5*(4), 281–300.

Bouveyron, C. and J. Jacques (2014). *funHDDC: Model-based clustering in group-specific functional subspaces*. R package version 1.0.

Cattell, R. B. (1966). The scree test for the number of factors. *Multivariate Behaviour Research* (2), 245–276.

Cuesta-Albertos, J. A. and R. Fraiman (2007). Impartial trimmed k-means for functional data. *Comput. Statist. Data Anal. 51*(10), 4864–4877.

Cuesta-Albertos, J. A., A. Gordaliza, and C. Matrán (1997). Trimmed k-means: an attempt to robustify quantizers. *Ann. Statist. 25*(2), 553–576.

Delaigle, A. and P. Hall (2010). Defining probability density for a distribution of random functions. *Ann. Statist. 38*(2), 1171–1193.

Febrero, M., P. Galeano, and W. González-Manteiga (2008). Outlier detection in functional data by depth measures, with application to identify abnormal NO_x levels. *Environmetrics 19*(4), 331–345.

Febrero-Bande, M. and M. Oviedo de la Fuente (2012). Statistical computing in functional data analysis: The R package fda.usc. *Journal of Statistical Software 51*(4), 1–28.

Ferraty, F. and P. Vieu (2006). *Nonparametric functional data analysis*. Springer Series in Statistics. Springer, New York.

Fraley, C. and A. E. Raftery (2002). Model-based clustering, discriminant analysis, and density estimation. *J. Amer. Statist. Assoc. 97*(458), 611–631.

Fritz, H., L. A. García-Escudero, and A. Mayo-Iscar (2013). A fast algorithm for robust constrained clustering. *Comput. Statist. Data Anal. 61*, 124–136.

Gallegos, M. T. (2002). Maximum likelihood clustering with outliers. In *Classification, clustering, and data analysis (Cracow, 2002)*, Stud. Classification Data Anal. Knowledge Organ., pp. 247–255. Springer, Berlin.

García-Escudero, L. A. and A. Gordaliza (2005). A proposal for robust curve clustering. *J. Classification 22*(2), 185–201.

García-Escudero, L. A., A. Gordaliza, C. Matrán, and A. Mayo-Iscar (2008). A general trimming approach to robust cluster analysis. *Ann. Statist. 36*(3), 1324–1345.

García-Escudero, L. A., A. Gordaliza, C. Matrán, and A. Mayo-Iscar (2015). Avoiding spurious local maximizers in mixture modeling. *Stat. Comput. 25*(3), 619–633.

García-Escudero, L. A., A. Gordaliza, and A. Mayo-Iscar (2014). A constrained robust proposal for mixture modeling avoiding spurious solutions. *Adv. Data Anal. Classif. 8*(1), 27–43.

Jacques, J. and C. Preda (2013). Funclust: A curves clustering method using functional random variables density approximation. *Neurocomputing 112*, 164–171.

James, G. M. and C. A. Sugar (2003). Clustering for sparsely sampled functional data. *J. Amer. Statist. Assoc. 98*(462), 397–408.

Ramsay, J. O. and B. W. Silverman (2005). *Functional data analysis* (Second ed.). Springer Series in Statistics. Springer, New York.

Ritter, G. (2015). *Robust cluster analysis and variable selection*, Volume 137 of *Monographs on Statistics and Applied Probability*. CRC Press, Boca Raton, FL.

Sawant, P., N. Billor, and H. Shin (2012). Functional outlier detection with robust functional principal component analysis. *Comput. Statist. 27*(1), 83–102.

Sguera, C., P. Galeano, and R. E. Lillo (2015, 7). Functional outlier detection by a local depth with application to nox levels. *Stochastic Environmental Research and Risk Assessment 28*(462), 1835–1851.

Soueidatt, M. (2014). *Funclustering: A package for functional data clustering.* R package version 1.0.1.